Generalized Tag-induced Cross-Domain Collaborative Filtering

YUE SHI, Delft University of Technology
MARTHA LARSON, Delft University of Technology
ALAN HANJALIC, Delft University of Technology

arXiv:1302.4888v1 [cs.IR] 20 Feb 2013

One of the most challenging problems in recommender systems based on the collaborative filtering (CF) concept is data sparseness, i.e., limited user preference data is available for making recommendations. Cross-domain collaborative filtering (CDCF) has been studied as an effective mechanism to alleviate data sparseness of one domain by transferring knowledge about user preferences from other domains. However, there are two key issues that need to be addressed to make a CDCF approach successful: (a) what common characteristics can be used to establish a link between different domains and (b) how to get each domain effectively and efficiently benefit from such a link. In this paper, we propose a novel algorithm, Generalized Tag-induced Cross-domain Collaborative Filtering (GTagCDCF), that exploits user-contributed tags as common characteristics to link different recommender domains together. Formulated from the probabilistic point of view, GTagCDCF takes into account all the user-item relations, the user-tag relations and the item-tag relations from different domains, resulting in a substantially extended collective matrix factorization framework, in which the common tags take the role of effectively transferring the knowledge between different domains. GTagCDCF is also found to be efficient, since its complexity is linear in the number of observed relations among users, items and tags from all domains. Using publicly available CF datasets to represent three cross-domain cases, i.e., two two-domain cases and one three-domain case, we experimentally demonstrate that GTagCDCF substantially outperforms several state-of-the-art single domain and cross-domain CF-based recommendation approaches. GTagCDCF is also shown to be effective for heterogeneous cross-domain cases, in which different domains are characterized by different types of user preferences. In addition, our investigation of the impact of user tagging behavior on GTagCDCF led to the conclusion that users can already benefit from GTagCDCF if they only share a few common tags. Finally, we validate the robustness of GTagCDCF with respect to the scale of datasets and the number of domains, based on a three-domain experiment.

Categories and Subject Descriptors: H.3.3 [**Information Storage and Retrieval**]: Information Search and Retrieval - Information Filtering

General Terms: Algorithms, Experimentation, Performance

Additional Key Words and Phrases: Collaborative filtering, cross domain collaborative filtering, matrix factorization, tag, recommender systems

ACM Reference Format:
Shi, Y., Larson, M., and Hanjalic, A. xxxx. Generalized tag-induced cross-domain collaborative filtering. ACM x, x, Article xx (March xxxx), 25 pages.
DOI = 10.1145/0000000.0000000 http://doi.acm.org/10.1145/0000000.0000000

Portions of this work are based on the paper "Tags as bridges between domains: Improving recommendation with tag-induced cross-domain collaborative filtering" [Shi et al. 2011], which appears in Proceedings of UMAP'11, Girona, Spain, July 2011.
Author's addresses: Y. Shi, M. Larson, A. Hanjalic, Multimedia Information Retrieval Lab, Department of Intelligent Systems, Faculty of Electrical Engineering, Mathematics and Computer Science, Delft University of Technology, Mekelweg 4, 2628CD Delft, The Netherlands; e-mail: {y.shi, m.a.larson, a.hanjalic}@tudelft.nl.
Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.
© xxxx ACM 0000-0000/xxxx/03-ARTxx $10.00
DOI 10.1145/0000000.0000000 http://doi.acm.org/10.1145/0000000.0000000

1. INTRODUCTION

Recommender systems have become ubiquitous to relieve internet users from information overload in online sites, such as Amazon[1], YouTube[2], Netflix[3]. One of the most successful techniques in recommender systems is collaborative filtering [Adomavicius and Tuzhilin 2005], which is based on the concept that users who have similar interest in the past may still have similar preference in future [Resnick et al. 1994]. However, one of the most challenging problems in CF has been recognized as data sparseness [Adomavicius and Tuzhilin 2005; Cacheda et al. 2011; Herlocker et al. 2004], since most of the users in a recommender system do not rate (or give any kind of feedback to) many items, resulting in highly sparse user-item relations. Specifically to address the data sparseness problem, recent research in recommender systems has started to study the potential of cross-domain collaborative filtering (CDCF) [Li et al. 2009a; Li et al. 2009b; Li et al. 2011; Pan et al. 2010; Pan et al. 2011; Zhang et al. 2010]. The key idea of CDCF is to make use of rating (or user preference) data from different recommender domains in order to benefit recommendations of one domain from other domains. Conceptually, CDCF attains alleviation of data sparseness in one recommender domain, by means of transferring knowledge of user preference from other domains.

Although the idea of CDCF is intuitively sound, there are several challenges that we have to face in order to develop a successful CDCF approach. First, different recommender domains are usually mutually exclusive. Each domain involves a certain type of products, e.g., movies, music or books, which could be different from products in other domains, and each domain involves a certain set of users, which could also be different from users in other domains or hardly distinguishable from their usernames. As a result, it is difficult, or perhaps also undesirable, to directly link users or items from different domains. Second, different recommender domains may probably have different types of user preference feedback. For instance, one domain may encode user preference explicitly, such as ratings or scores, while another domain may only have user preference with implicit feedback, such as the number of clicks, the duration of staying in a site, etc. As a result, it is difficult to utilize those heterogeneous user preference data from different domains for mutual benefits. Third, even if common characteristics among different domains are available, it is still difficult to maximally benefit from those common characteristics for improving recommendation quality in each domain, while keeping the recommendation process scalable.

In this paper, we propose a novel approach, generalized tag-induced cross-domain collaborative filtering (GTagCDCF), that specifically addresses the aforementioned challenges. We exploit user-generated tags that are common to different domains to link their users and items, in order to transfer the knowledge of user preference from one domain to another. Tagging has become a ubiquitous function in most of today's recommender systems [Robu et al. 2009; Song et al. 2011], in which a user can annotate an item with an arbitrary textual expression, i.e., a tag. Since a tag usually encodes a user's feeling about an item and even users from different domains may use same tags [Robu et al. 2009], we infer that users from different domains may use same tags to express similar feeling to different domain items. GTagCDCF is formulated from probabilistic point of view that takes into account all of the user-item preference, the user-tag relations and the item-tag relations from different domains, resulting in a matrix factorization framework, from which all the latent features of users and items from different domains can be learned with tag-induced benefits from each other. GTagCDCF is a generalized framework for CDCF, since 1) it maximally

[1]http://www.amazon.com/
[2]http://www.youtube.com/
[3]http://www.netflix.com/

exploits the information of common tags for knowledge transfer among domains, 2) it suits most of the cross-domain situations, i.e., being able to tackle different sets of users and items in different domains.

The novel contribution of this paper can be summarized as follows.

— We propose a new cross-domain collaborative filtering framework that could exploit explicitly shared knowledge between different domains for mutual improvement of recommendation performance.
— We experimentally demonstrate that the proposed GTagCDCF could significantly outperform several single-domain and cross-domain CF approaches.
— We find that the proposed GTagCDCF could particularly benefit for users who are active in tagging.

The remainder of the paper is structured as follows. In the next section, we summarize related work and position our approach with respect to it. We formalize the research problem and the terminology in section 3, followed by a detailed presentation of GTagCDCF in section 4. In section 5, we present experimental evaluation of GTagCDCF on three cross-domain cases. The last section sums up the key aspects of GTagCDCF and briefly addresses the direction for future work.

2. RELATED WORK

This section briefly summarizes the existing related research in CF, CDCF and tag-aware recommendation, in order to position the approach we propose in this paper.

2.1. Collaborative Filtering

CF approaches are usually categorized as either memory-based or model-based [Adomavicius and Tuzhilin 2005; Breese et al. 1998; Herlocker et al. 2004]. A recent comparative study of CF approaches can be found in [Cacheda et al. 2011]. Depending on whether the recommendation for a user is aggregated from some other users with similar interest, or from items that are similar to those she already liked, memory-based CF approaches can be further categorized as either user-based collaborative filtering (UBCF) [Herlocker et al. 1999; Resnick et al. 1994] or item-based collaborative filtering [Deshpande and Karypis 2004; Linden et al. 2003; Sarwar et al. 2001]. Although various modifications have been made for memory-based approaches, e.g., similarity fusion [Wang et al. 2006] and employing external knowledge [Umyarov and Tuzhilin 2011], the key drawback remains in the expensive computation for similarities among all users or items, which could hardly scale with the extremely large numbers of users and items in real-world recommender systems. Compared to memory-based approaches, model-based approaches, which do not rely on computing similarities among users or items, first use a training set of user-item preference data to learn a prediction model, then apply that model to generate recommendations. Conventional model-based CF approaches include Gaussian mixture model [Kleinberg and Sandler 2004; Si and Jin 2003] and latent semantic model [Hofmann 2004]. Recently, matrix factorization (MF) techniques have attracted much research attention, due to the advantages of scalability and accuracy, especially for large-scale data, as exemplified by the Netflix contest [Koren et al. 2009; Paterek 2007; Salakhutdinov and Mnih 2008]. Generally, MF techniques exploit the observed user preference data of a recommender system to learn both latent user features and latent item features, which are further used to predict unknown user preference to items. Probabilistic matrix factorization (PMF) [Salakhutdinov and Mnih 2008] illustrates the rating factorization from a probabilistic point of view. Our work in this paper builds on the PMF concept, but goes beyond PMF due to the substantial difference between cross-domain scenarios and single-domain scenarios.

We notice that our work is closely related to the work of MF with multiple matrices. To the best of our knowledge, under the context of CF, collective matrix factorization (CMF) [Singh and Gordon 2008] is the first work that proposed to factorize multiple matrices related to either users or items in addition to user-item rating matrix, while another similar work was proposed a little earlier under the context of text classification [Zhu et al. 2007]. The proposed GTagCDCF in this paper can be viewed as an extended version of CMF. The detailed analysis of the substantial difference between GTagCDCF and CMF will be presented in Section 4.5, after introducing the technical details of GTagCDCF. Recently, due to the availability of large amount data of social relationship and contextual information in recommender systems, researchers have proposed to improve recommendation by developing MF with multiple matrices involving not only user-item rating matrix but also social-aware information [Ma et al. 2011] or context-aware information [Shi et al. 2013]. In addition, researchers have also proposed to model MF with multiple matrices from probabilistic point of view by extending the PMF framework with additional priors [Adams et al. 2010; Porteous et al. 2010]. However, compared to the aforementioned work, our work in this paper is substantially different, since we tackle the CDCF problem, in which a key challenge lies in the connection between different domains.

2.2. Cross-domain Collaborative Filtering

We notice that the most closely related work to our work is on the CDCF approaches. One of the earliest work on CDCF was done by Berkovsky et al. [Berkovsky et al. 2007], who deployed several mediation approaches for importing and aggregating user rating vectors from different domains. Recently, the research of CDCF has been influenced and benefitted from the progress of transfer learning [Pan and Yang 2010], a machine learning paradigm for sharing knowledge among different domains. Coordinate system transfer [Pan et al. 2010] is proposed to first learn latent features of users and items from an auxiliary domain (which has relatively more user preference data), and then adapt them to a target domain (which has relatively less user preference data) in order to regularize the learning of latent features for the users and the items in the target domain. Further extension has also been proposed to exploit implicit user feedbacks rather than explicit user ratings to constitute the auxiliary domain [Pan et al. 2011]. However, these approaches require that either users or items are shared between the domains, which is, as already mentioned in Section 1, a condition not commonly encountered in practical applications. Codebook transfer (CBT) [Li et al. 2009a] and rating-matrix generative model (RMGM) [Li et al. 2009b] are proposed to transfer knowledge from an auxiliary domain by learning an implicit cluster-level rating pattern that could be shared to a target domain. Similarly, multi-domain CF is proposed to extend PMF in multiple domains involving explicit user preference [Zhang et al. 2010] or implicit user feedbacks [Tang et al. 2011], by means of learning an implicit correlation matrix, which is supposed to link different domains for knowledge transfer. One of the latest work has adopted the CDCF framework of RMGM [Li et al. 2009b] to address the problem of dynamic CF [Li et al. 2011]. Compared to all the aforementioned CDCF approaches, the proposed GTagCDCF in this paper is substantially different in that we exploit an explicit information source, i.e., the tag, between different domains for knowledge transfer, rather than relying on implicit domain correlations that are still mined from user preference data. In addition, GTagCDCF is different from CBT and RMGM, since GTagCDCF could tackle the case of domains with different types of user preference, while CBT and RMGM require that different domains should have the same rating scale.

In our previous work, we have proposed tag-induced cross-domain collaborative filtering (TagCDCF) [Shi et al. 2011] to use common tags as bridges to link different

domains for improving CDCF. However, there are two specific drawbacks that needs to be addressed. First, TagCDCF links different domains based on the cross-domain similarities, which need to be computed by using a specific similarity function. However, the choice of the similarity function might be sensitive to the recommendation performance, and might vary substantially for different user cases. For this reason, a good choice from a number of similarity functions usually requires costly empirical investigation. In addition, calculating these similarities offline is still a computationally expensive task in the case of large number of cross-domain users or items. Second, TagCDCF only exploits user-tag relations and item-tag relations with binary indicators, while the frequency of a user assigning a tag and the frequency of an item annotated by a tag are totally ignored. As a result, valuable information may lose during this dramatic simplification, leading to a suboptimal knowledge transfer by the common tags. The proposed GTagCDCF specifically addresses the two drawbacks of TagCDCF, as it does not require computing cross-domain similarities and does explicitly take into account the frequency of user-tag relations and item-tag relations. For this reason, we refer to GTagCDCF as a generalized version of TagCDCF, since it suits more general use cases.

2.3. Tag-aware Recommendation

We also point out that our work in this paper is related to tag-aware recommendation. Tags are exploited to improve recommendation by various means, e.g., to fuse tags into traditional user-based CF and item-based CF [Liang et al. 2010; Sen et al. 2009; Tso-Sutter et al. 2008], to incorporate tags into probabilistic latent semantic model [Hofmann 2004] to unify user-item relations and item-tag relations [Wang et al. 2006; Wetzker et al. 2009], and to use tag-based user correlations as a regularization for PMF [Zhen et al. 2009], etc. More recently, another group of state-of-the-art approaches has employed tensor factorization techniques [Kolda and Sun 2008] for tag-aware recommendation. Under such approaches, item recommendations or tag recommendations are learned from the {user, tag, item} triplet/ternary data directly [Rendle and Schmidt-Thieme 2010; Symeonidis et al. 2010]. Given the advantages of using tags for improving recommender systems, our work in this paper goes a step further that explores the potential of using tags for introducing mutual benefits between different recommender domains.

3. PROBLEM STATEMENT AND TERMINOLOGY

3.1. Problem Statement

The research problem we study in this paper can be stated as: *How to effectively exploit the common tags between different recommender domains for improving the quality of recommendations in each domain.* Specifically, we illustrate this problem by a toy example as shown in Fig. 1. Supposing that Alice is a user in a movie recommender system/domain, and Bob is a user in a book recommender domain, the problem is to predict Alice's rating on Movie2 and Bob's rating on Book1 (or in other words, whether the movie recommender system should recommend Movie2 to Alice and whether the book recommender system should recommend Book1 to Bob). Note that in the example we use a solid line to denote a rating (e.g., scaled from 1 to 5), and a dashed line to denote a relationship between a user and a tag or a relationship between an item and a tag. For this example, we can see that the two domains may both suffer from data sparseness, since it is difficult to infer the user's preference solely based on the domain that the user exists in. However, there is potential to infer the user's preference based on the common tags "fun" and "ridiculous". As can be seen, Alice rated the Movie1 with the highest rating, and she also tagged Movie1 with "fun". From this observation,

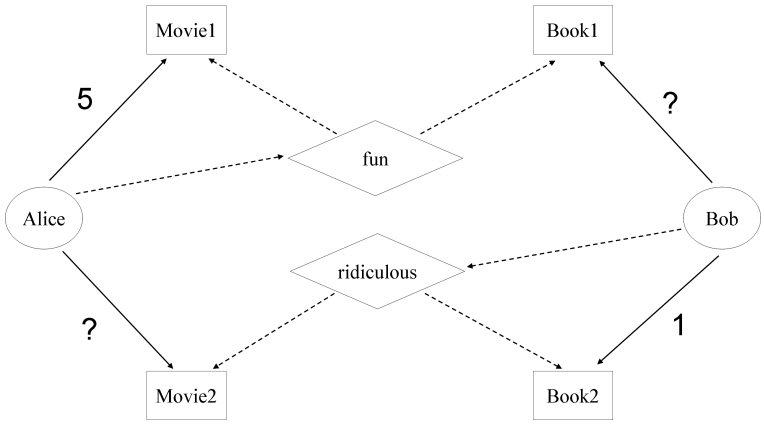


Fig. 1. A toy example of two recommender domains with common tags

we may infer that an item tagged by "fun" tend to be favored by users. Based on this inference, we can infer that in the book domain, Bob might also be in favor of Book1, which is tagged with "fun" by some other users. According to the same reasoning, we can infer that Movie2 may not be a good recommendation for Alice. Although this toy example in Fig. 1 is much simplified from real systems, it still demonstrates that there is great potential of mutually benefitting different recommender domains from the common tags. Our work in this paper is to develop a new algorithm that can effectively achieve the aim of cross-domain CF via common tags.

3.2. Terminology and Notations

In this subsection, we first define the terms used further in the paper. A user-item matrix in kth ($k = 1, 2, \ldots, K$) domain is denoted as $\mathbf{R}^{(\mathbf{k})}$, which consists of M_k users and N_k items. Note that in this paper, $\mathbf{R}^{(\mathbf{k})}$ could be a matrix with either explicit ratings or implicit feedback data. We use $R_{ij}^{(k)}$ to denote user i's preference score (derived from either explicit rating or implicit feedback) to item j in the kth domain. A user-tag matrix in kth domain is denoted as $\mathbf{F}^{(\mathbf{U_k})}$, in which $F_{il}^{(U_k)}$ indicates the frequency of user i using tag l. Similarly, an item-tag matrix in kth domain is denoted as $\mathbf{F}^{(\mathbf{V_k})}$, in which $F_{jl}^{(V_k)}$ indicates the frequency of item j annotated by tag l. Note that without loss of generality, the non-zero entries of $\mathbf{R}^{(\mathbf{k})}$, $\mathbf{F}^{(\mathbf{U_k})}$ and $\mathbf{F}^{(\mathbf{V_k})}$ are, respectively, normalized to be within the range [0,1], by means of dividing over the maximal value in each matrix. We adopt the convention of denoting the number of non-zero entries in a matrix $\mathbf{A}$ as $|\mathbf{A}|$. In addition, we use I^A to serve an indicator function, which gives $I_{ij}^A = 1$ if $A_{ij} > 0$, 0 otherwise. $\|\mathbf{A}\|_{Fro}$ denotes the Frobenius norm of matrix $\mathbf{A}$.

4. GENERALIZED TAG-INDUCED CROSS-DOMAIN COLLABORATIVE FILTERING

In this section, we present the detail of generalized tag-induced cross-domain collaborative filtering (GTagCDCF). We first introduce the graphical model of GTagCDCF, from which we demonstrate how tags bridge different domains. Then, we interpret the graphical model from probabilistic point of view, resulting in a matrix factorization framework that incorporates tags to benefit different domains in learning latent features of users and items. Finally, we present complexity analysis of GTagCDCF and summarize the main characteristics of GTagCDCF.

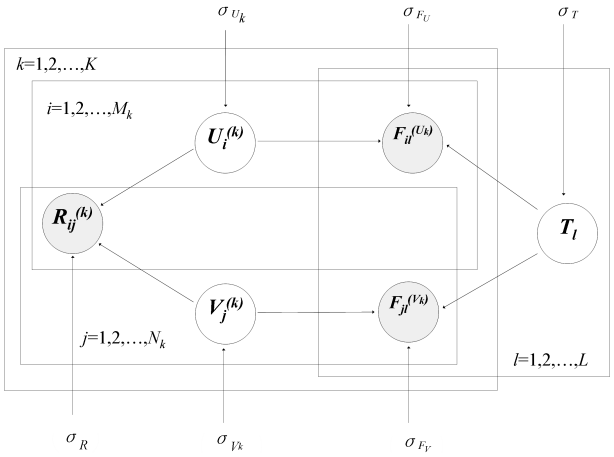


Fig. 2. Graphical model of GTagCDCF

4.1. Graphical Model

We present the graphical model of the proposed GTagCDCF approach, as shown in Fig. 2, in order to illustrate the relationships between different variables/factors from K domains. $\mathbf{U^{(k)}}$ denotes a $d \times M_k$ matrix, whose ith column, i.e., $U_i^{(k)}$, represents a d-dimensional latent feature vector of user i in the k domain. $\mathbf{V^{(k)}}$ denotes a $d \times N_k$ matrix, whose jth column i.e., $V_j^{(k)}$, represents a d-dimensional latent feature vector of user j in the kth domain. Supposing there are L tags that are common (co-existing) to all the K domains, $\mathbf{T}$ denotes a $d \times L$ matrix, whose lth column, i.e., T_l represents a d-dimensional latent feature vector of tag l. Our objective is to estimate the unknown latent features of users and items, i.e., $\mathbf{U^{(k)}}$ and $\mathbf{V^{(k)}}$ (k=1,2…,K), which can be then used to predict recommendations in each domain. Compared to PMF [Salakhutdinov and Mnih 2008], we can observe that in Fig. 2, the sub-graph that only involves $\mathbf{U^{(k)}}$, $\mathbf{V^{(k)}}$ and $\mathbf{R^{(k)}}$ is equivalent to PMF in a single domain. The set of latent features of common tags, i.e., $\mathbf{T}$, takes role of the central node that connects K domains. As can be seen, the latent features of users in one domain would influence those in other domains via the common tags, and so do the latent features of items. The strength of the mutual influence is reflected by the frequency of user-tag relations, i.e., $\mathbf{F^{(U_k)}}$ and the frequency of item-tag relations, i.e., $\mathbf{F^{(V_k)}}$. In the following, we will present the detail of learning latent features of users and items based on this graphical model.

4.2. Model Realization and Optimization

In this subsection, we interpret the graphical model of GTagCDCF in Fig. 2 from a probabilistic point of view, and further derive the objective function of GTagCDCF. For notation convenience, we introduce some substitutions for the variables in the graphical model, as shown below:

$$\mathbf{Y} = \{\mathbf{R^{(k)}}, \mathbf{F^{(U_k)}}, \mathbf{F^{(V_k)}} | k = 1, 2, \ldots, K\}$$
$$\mathbf{X} = \{\mathbf{U^{(k)}}, \mathbf{V^{(k)}}, \mathbf{T} | k = 1, 2, \ldots, K\}$$
$$\mathbf{\Theta} = \{\sigma_R, \sigma_{F_U}, \sigma_{F_V}, \sigma_{U_k}, \sigma_{V_k}, \sigma_T | k = 1, 2, \ldots, K\}$$

According to the probabilistic graphical model theory [Bishop 2006], the joint probability distribution of all the variables in the graphical model can be factorized as below:

$$
\begin{aligned}
& p(\mathbf{X}, \mathbf{Y}, \mathbf{\Theta}) \qquad (1)\\
&= \left(\prod_{k=1}^{K} p(\mathbf{R^{(k)}} | \mathbf{U^{(k)}}, \mathbf{V^{(k)}}, \sigma_R) p(\mathbf{F^{(U_k)}} | \mathbf{U^{(k)}}, \mathbf{T}, \sigma_{F_U}) p(\mathbf{F^{(V_k)}} | \mathbf{V^{(k)}}, \mathbf{T}, \sigma_{F_V}) p(\mathbf{U^{(k)}} | \sigma_{U_k}) p(\mathbf{V^{(k)}} | \sigma_{V_k}) \right) \\
& p(\mathbf{T} | \sigma_T) p(\mathbf{\Theta})
\end{aligned}
$$

Applying product rule to the left side of Eq. (1), i.e., $p(\mathbf{X}, \mathbf{Y}, \mathbf{\Theta}) = p(\mathbf{X} | \mathbf{Y}, \mathbf{\Theta}) p(\mathbf{Y}, \mathbf{\Theta})$, and neglecting constant prior probabilities, i.e., priors that have no influence on latent features $\mathbf{U^{(1)}}$, $\mathbf{V^{(1)}}$, $\mathbf{U^{(2)}}$, $\mathbf{V^{(2)}}$, and $\mathbf{T}$, we obtain the conditional probability distribution of latent features as below:

$$
\begin{aligned}
& p(\mathbf{X} | \mathbf{Y}, \mathbf{\Theta}) \qquad (2)\\
&\propto \left(\prod_{k=1}^{K} p(\mathbf{R^{(k)}} | \mathbf{U^{(k)}}, \mathbf{V^{(k)}}, \sigma_R) p(\mathbf{F^{(U_k)}} | \mathbf{U^{(k)}}, \mathbf{T}, \sigma_{F_U}) p(\mathbf{F^{(V_k)}} | \mathbf{V^{(k)}}, \mathbf{T}, \sigma_{F_V}) p(\mathbf{U^{(k)}} | \sigma_{U_k}) p(\mathbf{V^{(k)}} | \sigma_{V_k}) \right) \\
& p(\mathbf{T} | \sigma_T)
\end{aligned}
$$

As can be seen in Eq. (2), the latent features of users, items and tags can be estimated by the principle of maximum a posteriori, if we interpret all the conditional probabilities in the right side of Eq. (2) properly. In the following, we will present interpretations of all the conditional probabilities based on insights from probabilistic matrix factorization [Salakhutdinov and Mnih 2008].

By using the conditional independence of individual ratings given latent features of users and items, the conditional probability distribution over observed ratings in the kth domain is defined as:

$$
p(\mathbf{R^{(k)}} | \mathbf{U^{(k)}}, \mathbf{V^{(k)}}, \sigma_R) = \prod_{i=1}^{M_k} \prod_{j=1}^{N_k} \left[\mathcal{N} \left(R_{ij}^{(k)} | g(U_i^{(k)T} V_j^{(k)}), \sigma_R^2 \right) \right]^{I_{ij}^{R^{(k)}}} \qquad (3)
$$

where $\mathcal{N}(x | \mu, \sigma^2)$ denotes the probability density function for a Gaussian distribution with mean value μ and variance σ^2. $g(x)$ is a logistic function, i.e., $g(x) = 1/(1 + e^{-x})$, which is used to map the inner product of latent user features and latent item features to be within the range [0,1].

Similarly, we also define the conditional probability distributions over observed user-tag frequency and observed item-tag frequency in the kth domain as in Eq. (4) and in Eq. (5), respectively.

$$
p(\mathbf{F^{(U_k)}} | \mathbf{U^{(k)}}, \mathbf{T}, \sigma_{F_U}) = \prod_{i=1}^{M_k} \prod_{l=1}^{L} \left[\mathcal{N} \left(F_{il}^{(U_k)} | g(U_i^{(k)T} T_l), \sigma_{F_U}^2 \right) \right]^{I_{il}^{F^{(U_k)}}} \qquad (4)
$$

$$
p(\mathbf{F^{(V_k)}} | \mathbf{V^{(k)}}, \mathbf{T}, \sigma_{F_V}) = \prod_{j=1}^{N_k} \prod_{l=1}^{L} \left[\mathcal{N} \left(F_{jl}^{(V_k)} | g(V_j^{(k)T} T_l), \sigma_{F_V}^2 \right) \right]^{I_{jl}^{F^{(V_k)}}} \qquad (5)
$$

Finally, we represent latent features of users and items in each domain, and latent features of tags with zero-mean spherical Gaussian priors [Tipping and Bishop 1999].

$$p(\mathbf{U^{(k)}}|\sigma_{U_k}) = \prod_{i=1}^{M_k} \mathcal{N}\left(U_i^{(k)}|0, \sigma_{U_k}^2 \mathbf{\Lambda}\right) \tag{6}$$

$$p(\mathbf{V^{(k)}}|\sigma_{V_k}) = \prod_{j=1}^{N_k} \mathcal{N}\left(V_j^{(k)}|0, \sigma_{V_k}^2 \mathbf{\Lambda}\right) \tag{7}$$

$$p(\mathbf{T}|\sigma_T) = \prod_{l=1}^{L} \mathcal{N}\left(T_l|0, \sigma_T^2 \mathbf{\Lambda}\right) \tag{8}$$

where $\mathbf{\Lambda}$ is a $d \times d$ identity matrix.

Substituting Eq. (3$\sim$8) into Eq. (2), we obtain the posterior as below:

$$\begin{aligned} & p(\mathbf{X}|\mathbf{Y}, \mathbf{\Theta}) \\ & \propto \prod_{k=1}^{K} \prod_{i=1}^{M_k} \prod_{j=1}^{N_k} \left[\mathcal{N}\left(R_{ij}^{(k)}|g(U_i^{(k)T} V_j^{(k)}), \sigma_R^2\right)\right]^{I_{ij}^{R^{(k)}}} \prod_{i=1}^{M_k} \prod_{l=1}^{L} \left[\mathcal{N}\left(F_{il}^{(U_k)}|g(U_i^{(k)T} T_l), \sigma_{F_U}^2\right)\right]^{I_{il}^{F^{(U_k)}}} \\ & \prod_{j=1}^{N_k} \prod_{l=1}^{L} \left[\mathcal{N}\left(F_{jl}^{(V_k)}|g(V_j^{(k)T} T_l), \sigma_{F_V}^2\right)\right]^{I_{jl}^{F^{(V_k)}}} \prod_{i=1}^{M_k} \mathcal{N}\left(U_i^{(k)}|0, \sigma_{U_k}^2 \mathbf{\Lambda}\right) \prod_{j=1}^{N_k} \mathcal{N}\left(V_j^{(k)}|0, \sigma_{V_k}^2 \mathbf{\Lambda}\right) \\ & \prod_{l=1}^{L} \mathcal{N}\left(T_l|0, \sigma_T^2 \mathbf{\Lambda}\right) \end{aligned} \tag{9}$$

As mentioned before, the latent features of users, items and tags can be estimated by maximizing this posterior as in Eq. (9). Notice that it is also equivalent to achieve the maximization of the posterior by minimizing the negative log-posterior, which is shown below:

$$\begin{aligned} & -\log p(\mathbf{X}|\mathbf{Y}, \mathbf{\Theta}) \\ & \propto \sum_{k=1}^{K} \sum_{i=1}^{M_k} \sum_{j=1}^{N_k} \left[\frac{1}{2\sigma_R^2} I_{ij}^{R^{(k)}} \left(R_{ij}^{(k)} - g(U_i^{(k)T} V_j^{(k)})\right)^2\right] + \sum_{k=1}^{K} \sum_{i=1}^{M_k} \sum_{l=1}^{L} \left[\frac{1}{2\sigma_{F_U}^2} I_{il}^{F^{(U_k)}} \left(F_{il}^{(U_k)} - g(U_i^{(k)T} T_l)\right)^2\right] \\ & + \sum_{k=1}^{K} \sum_{j=1}^{N_k} \sum_{l=1}^{L} \left[\frac{1}{2\sigma_{F_V}^2} I_{jl}^{F^{(V_k)}} \left(F_{jl}^{(V_k)} - g(V_j^{(k)T} T_l)\right)^2\right] + \sum_{k=1}^{K} \sum_{i=1}^{M_k} \frac{1}{2\sigma_{U_k}^2} \left(U_i^{(k)T} U_i^{(k)}\right) \\ & + \sum_{k=1}^{K} \sum_{j=1}^{N_k} \frac{1}{2\sigma_{V_k}^2} \left(V_j^{(k)T} V_j^{(k)}\right) + \sum_{l=1}^{L} \frac{1}{2\sigma_T^2} \left(T_l^T T_l\right) + C \end{aligned} \tag{10}$$

in which

$$\begin{aligned} C = & \sum_{k=1}^{K} \sum_{i=1}^{M_k} \sum_{j=1}^{N_k} I_{ij}^{R^{(k)}} \log \sqrt{2\pi}\sigma_R + \sum_{k=1}^{K} \sum_{i=1}^{M_k} \sum_{l=1}^{L} I_{il}^{F^{(U_k)}} \log \sqrt{2\pi}\sigma_{F_U} + \sum_{k=1}^{K} \sum_{j=1}^{N_k} \sum_{l=1}^{L} I_{il}^{F^{(V_k)}} \log \sqrt{2\pi}\sigma_{F_V} \\ & + \sum_{k=1}^{K} dM_k \log \sqrt{2\pi}\sigma_{U_k} + \sum_{k=1}^{K} dN_k \log \sqrt{2\pi}\sigma_{V_k} + dL \log \sqrt{2\pi}\sigma_T \end{aligned} \tag{11}$$

Note that C is irrelevant to latent features $\mathbf{U^{(k)}}$, $\mathbf{V^{(k)}}$, and $\mathbf{T}$, meaning it has no influence on optimizing these latent features towards an optimal solution. In order to further simplify the model, we also assume that prior variance of all the latent features are the same, i.e., $\sigma_{U_k}^2 = \sigma_{V_k}^2 = \sigma_T^2 = c$. Taking the above conditions into Eq. (10), we obtain

the final objective function of GTagCDCF as $G(\mathbf{X})$:

$$
\begin{aligned}
& G(\mathbf{X}) \qquad (12)\\
=& \frac{1}{2}\sum_{k=1}^{K}\sum_{i=1}^{M_k}\sum_{j=1}^{N_k}\left[I_{ij}^{R^{(k)}}\left(R_{ij}^{(k)}-g(U_i^{(k)T}V_j^{(k)})\right)^2\right]+\frac{\alpha}{2}\sum_{k=1}^{K}\sum_{i=1}^{M_k}\sum_{l=1}^{L}\left[I_{il}^{F^{(U_k)}}\left(F_{il}^{(U_k)}-g(U_i^{(k)T}T_l)\right)^2\right]\\
&+\frac{\beta}{2}\sum_{k=1}^{K}\sum_{j=1}^{N_k}\sum_{l=1}^{L}\left[I_{jl}^{F^{(V_k)}}\left(F_{jl}^{(k)}-g(V_j^{(k)T}T_l)\right)^2\right]+\frac{\lambda}{2}\left[\sum_{k=1}^{K}(\|\mathbf{U^{(k)}}\|_{Fro}^2+\|\mathbf{V^{(k)}}\|_{Fro}^2)+\|\mathbf{T}\|_{Fro}^2\right]
\end{aligned}
$$

where we have set substitutions as $\alpha = \sigma_R^2/\sigma_{F_U}^2$, $\beta = \sigma_R^2/\sigma_{F_V}^2$, and $\lambda = \sigma_R^2/\sigma_{U_k}^2$. λ is a regularization parameter that penalizes the magnitude of latent features in order to alleviate over-fitting. α and β are regarded as tradeoff parameters, which control the relative influence from the user-tag matrices in K domains and the item-tag matrices in K domains, respectively. We emphasize that in the objective function the second term reflects that latent features of users from K domains are associated through latent features of shared tags, and the third term reflects that latent features of items from K domains are also associated through latent features of shared tags. Therefore, the proposed GTagCDCF bridges different domains by exploiting the common tags that have relationships with both users and items from different domains.

The objective function in Eq. (12) is not jointly convex to all the variables of latent features, i.e., $\mathbf{U^{(1)}}$, $\mathbf{V^{(1)}}$, $\mathbf{U^{(2)}}$, $\mathbf{V^{(2)}}$, and $\mathbf{T}$. We choose to apply gradient descend with respect to one of these variables alternatively and keep all the other variables fixed, in order to attain a local minimum solution for the objective function. Specifically, the gradients with respect to each variable can be computed as below:

$$
\begin{aligned}
\frac{\partial G}{\partial U_i^{(k)}} &= \sum_{j=1}^{N_k}\left[I_{ij}^{R^{(k)}}\left(g(U_i^{(k)T}V_j^{(k)})-R_{ij}^{(k)}\right)g'(U_i^{(k)T}V_j^{(k)})V_j^{(k)}\right] \qquad (13)\\
&+\alpha\sum_{l=1}^{L}\left[I_{il}^{F^{(U_k)}}\left(g(U_i^{(k)T}T_l)-F_{il}^{(U_k)}\right)g'(U_i^{(k)T}T_l)T_l\right]+\lambda U_i^{(k)} \quad k=1,2\ldots,K\\
\frac{\partial G}{\partial V_j^{(k)}} &= \sum_{i=1}^{M_k}\left[I_{ij}^{R^{(k)}}\left(g(U_i^{(k)T}V_j^{(k)})-R_{ij}^{(k)}\right)g'(U_i^{(k)T}V_j^{(k)})U_i^{(k)}\right] \qquad (14)\\
&+\beta\sum_{l=1}^{L}\left[I_{jl}^{F^{(V_k)}}\left(g(V_j^{(k)T}T_l)-F_{jl}^{(V_k)}\right)g'(V_j^{(k)T}T_l)T_l\right]+\lambda V_j^{(k)} \quad k=1,2\ldots,K\\
\frac{\partial G}{\partial T_l} &= \alpha\sum_{k=1}^{K}\sum_{i=1}^{M_k}\left[I_{il}^{F^{(U_k)}}\left(g(U_i^{(k)T}T_l)-F_{il}^{(U_k)}\right)g'(U_i^{(k)T}T_l)U_i^{(k)}\right] \qquad (15)\\
&+\beta\sum_{k=1}^{K}\sum_{j=1}^{N_k}\left[I_{jl}^{F^{(V_k)}}\left(g(V_j^{(k)T}T_l)-F_{jl}^{(V_k)}\right)g'(V_j^{(k)T}T_l)V_j^{(k)}\right]+\lambda T_l
\end{aligned}
$$

For the consideration of better readability, we summarize the algorithm of GTagCDCF in Algorithm 1.

4.3. Complexity Analysis

The complexity of the objective function of GTagCDCF in Eq. (12) is $O(d\sum_k(|\mathbf{R^{(k)}}|+|\mathbf{F^{(U_k)}}|+|\mathbf{F^{(V_k)}}|+M_k+N_k)+dL)$. The complexity of computing each gradient in Eq. (13∼15) is $O(d\sum_k(|\mathbf{R^{(k)}}|+|\mathbf{F^{(U_k)}}|+M_k))$, $O(d\sum_k(|\mathbf{R^{(k)}}|+|\mathbf{F^{(V_k)}}|+N_k))$, and $O(d\sum_k(|\mathbf{F^{(U_k)}}|+|\mathbf{F^{(V_k)}}|)+dL)$, respectively. Considering that in reality we usually have conditions as $|\mathbf{R_{(k)}}|>> M_k, N_k$, $|\mathbf{F^{(U_k)}}|>> M_k, L$, and $|\mathbf{F^{(V_k)}}|>> N_k, L$, the total complexity of GTagCDCF is in the order of $O(\sum_k(|\mathbf{R^{(k)}}|+|\mathbf{F^{(U_k)}}|+|\mathbf{F^{(V_k)}}|))$, which is linear in the

ALGORITHM 1: GTagCDCF

Input: Normalized user-item preference matrix $\mathbf{R^{(k)}}$ (k=1,2...,K), normalized user-tag frequency matrix $\mathbf{F^{(U_k)}}$ (k=1,2...,K), normalized item-tag frequency matrix $\mathbf{F^{(V_k)}}$ (k=1,2...,K), tradeoff parameters α, β, regularization parameter λ, stop condition ϵ.
Output: Predicted user-item preference matrix $\hat{\mathbf{R}}^{\mathbf{(k)}}$ (k=1,2...,K).
Initialize $\mathbf{U^{(k)[0]}}$, $\mathbf{V^{(k)[0]}}$ (k=1,2...,K) and $\mathbf{T^{[0]}}$ with random values;
$t = 0$;
Compute $G^{[t]}$ as in Eq. (12);
repeat
 $\eta = 1$;
 Compute $\frac{\partial G}{\partial U^{(k)[t]}}$, $\frac{\partial G}{\partial V^{(k)[t]}}$ and $\frac{\partial G}{\partial T^{[t]}}$ as in Eq. (13∼15) ;
 repeat
 $\eta = \eta/2$; // maximize learning step size
 until $G(\mathbf{U^{(k)[t]}} - \eta\frac{\partial G}{\partial U^{(k)[t]}}, \mathbf{V^{(k)[t]}} - \eta\frac{\partial G}{\partial V^{(k)[t]}}, \mathbf{T^{[t]}} - \eta\frac{\partial G}{\partial T^{[t]}}) < G^{[t]}$;
 $\mathbf{U^{(k)[t+1]}} = \mathbf{U^{(k)[t]}} - \eta\frac{\partial G}{\partial U^{(k)[t]}}$;
 $\mathbf{V^{(k)[t+1]}} = \mathbf{V^{(k)[t]}} - \eta\frac{\partial G}{\partial V^{(k)[t]}}$;
 $\mathbf{T^{[t+1]}} = \mathbf{T^{[t]}} - \eta\frac{\partial G}{\partial T^{[t]}}$;
 Compute $G^{[t+1]}$ as in Eq. (12);
 $f = 1 - G^{[t+1]}/G^{[t]}$;
 $t = t + 1$;
until $f \leq \epsilon$;
$\hat{\mathbf{R}}^{\mathbf{(k)}} = \mathbf{U^{(k)[t]^T}V^{(k)[t]}}$ (k=1,2...,K);

total number of the known preference scores in $\mathbf{R^{(k)}}$, the known user-tag relationships in $\mathbf{F^{(U_k)}}$, and the known item-tag relationships in $\mathbf{F^{(V_k)}}$ from all the domains. This analysis indicates that GTagCDCF is appropriate for large scale use cases.

4.4. A Compact Formulation

It is worth to mention that we can present the objective function of GTagCDCF in a compact form, which is mathematically equivalent to Eq. (12). The compact formulation is expressed below:

$$\begin{aligned} G(\mathbf{X}) =& \frac{1}{2}\sum_{i=1}^{M_0}\sum_{j=1}^{N_0}\left[I_{ij}^{\mathbb{R}}\left(\mathbb{R}_{ij} - g(\mathbb{U}_i^T\mathbb{V}_j)\right)^2\right] + \frac{\alpha}{2}\sum_{i=1}^{M_0}\sum_{l=1}^{L}\left[I_{il}^{\mathbb{F}^{(\mathbb{U})}}\left(\mathbb{F}_{il}^{(\mathbb{U})} - g(\mathbb{U}_i^T T_l)\right)^2\right] \\ &+ \frac{\beta}{2}\sum_{j=1}^{N_0}\sum_{l=1}^{L}\left[I_{jl}^{\mathbb{F}^{(\mathbb{V})}}\left(\mathbb{F}_{jl}^{(\mathbb{V})} - g(\mathbb{V}_j^T T_l)\right)^2\right] + \frac{\lambda}{2}\left[(\|\mathbb{U}\|_{Fro}^2 + \|\mathbb{V}\|_{Fro}^2) + \|\mathbf{T}\|_{Fro}^2\right] \end{aligned} \tag{16}$$

where $M_0 = M_1 + M_2 + \ldots + M_K$, and $N_0 = N_1 + N_2 + \ldots + N_K$. $\mathbb{R}$ is a block diagonal matrix that contains all the preference data from K domains; $\mathbb{F}^{(\mathbb{U})}$ is a stacked matrix that contains the user-tag relations from all the users of K domains; and $\mathbb{F}^{(\mathbb{V})}$ is a stacked matrix that contains the item-tag relations from all the items of K domains. Specifically, these variables are shown below:

$$\mathbb{R} = \begin{bmatrix} \mathbf{R^{(1)}} & & & \\ & \mathbf{R^{(2)}} & & \\ & & \cdots & \\ & & & \mathbf{R^{(K)}} \end{bmatrix}, \quad \mathbb{F}^{(\mathbb{U})} = \begin{bmatrix} \mathbf{F^{(U_1)}} \\ \mathbf{F^{(U_2)}} \\ \cdots \\ \mathbf{F^{(U_K)}} \end{bmatrix}, \quad \mathbb{F}^{(\mathbb{V})} = \begin{bmatrix} \mathbf{F^{(V_1)}} \\ \mathbf{F^{(V_2)}} \\ \cdots \\ \mathbf{F^{(V_K)}} \end{bmatrix} \tag{17}$$

Correspondingly, $\mathbb{U}$ and $\mathbb{V}$ are stacked matrices that contain the latent factors of users and items, respectively, from K domains, as shown below:

$$\mathbb{U} = \begin{bmatrix} \mathbf{U^{(1)}} \\ \mathbf{U^{(2)}} \\ \cdots \\ \mathbf{U^{(K)}} \end{bmatrix}, \quad \mathbb{V} = \begin{bmatrix} \mathbf{V^{(1)}} \\ \mathbf{V^{(2)}} \\ \cdots \\ \mathbf{V^{(K)}} \end{bmatrix} \tag{18}$$

We consider that the compact formulation of GTagCDCF could help practitioners to implement the algorithm more easily.

4.5. Discussion

We finalized our presentation of the proposed GTagCDCF approach with a brief summary and discussion of the ways in which the algorithm addresses the specific challenges faced in CDCF.

(1) **Explicit Knowledge Transfer.** GTagCDCF utilizes explicitly shared knowledge, in the form of common tags, to introduce mutual benefit between domains. Compared to related work [Li et al. 2009a; Li et al. 2009b; Pan et al. 2010; Pan et al. 2011], in which implicit correlations are mined from user preference data, we emphasize that the characteristic of explicit knowledge transfer could make GTagCDCF more effective in connecting users and items from different domains, leading to an improved recommendation performance in each domain.

(2) **Heterogeneous Knowledge Transfer.** GTagCDCF is suitable for cross-domain cases with heterogeneous user preference data. GTagCDCF can be widely applied in practice, since most recommender systems may have different mechanisms to allow users expressing their preference to items. Compared to the related work [Li et al. 2009a; Li et al. 2009b], in which the CDCF approaches can only tackle different domains with the same rating scale, we emphasize that GTagCDCF could enjoy the advantage of heterogeneous knowledge transfer, since the latent features of common tags provide a basis for learning the latent features of users and items with different types of preference data from different domains.

(3) **Generalized Knowledge Transfer.** Compared to our previous work [Shi et al. 2011], GTagCDCF is a generalized version of TagCDCF in that it directly exploits user-tag relations and item-tag relations, while not relying on any particular similarity measurement to capture the relations between users or items from different domains. In addition, we emphasize that GTagCDCF provides a generalized framework to utilize any explicit common characteristics (not only common tags) for introducing mutual benefits between recommender domains.

The ability of GTagCDCF to face these challenges is gratifying, but not entirely surprising, given the success of single-domain approaches to PMF that exploit the simultaneous factorization of multiple matrices. As mentioned in Section 2, collective matrix factorization (CMF) [Singh and Gordon 2008] is such an approach, proposed to factorize multiple matrices of metadata/attributes related to users and items. For a better understanding of the difference between GTagCDCF and CMF, we illustrate the graphical model of CMF under the same variable representations as used in Fig. 3. Compared with Fig. 2, it can be seen that the original CMF is a single-domain model, which learns the latent factors based on the rating data and the metadata (comparable to the tags in this paper) within one domain. In contrast, GTagCDCF is a cross-domain model, in which the latent factors from one domain could influence those in another domain based on the shared metadata (the common tags in this paper). Note that in case it would be appropriate to conflate individual user/item sets into one large set, cross-domain recommendation can also be expressed in the CMF framework. More

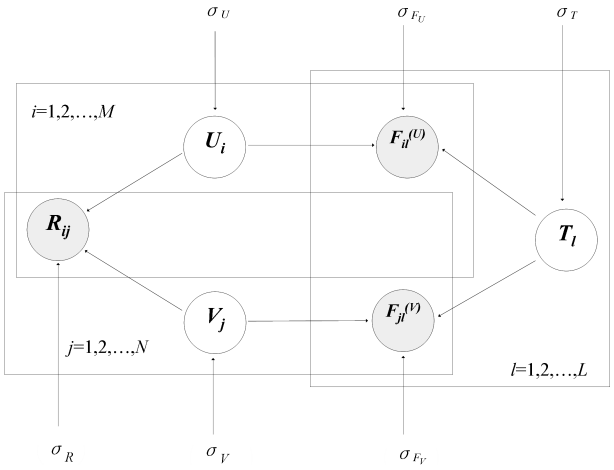


Fig. 3. Graphical model of CMF

deeply, CMF is essentially an extended version of PMF [Adams et al. 2010; Porteous et al. 2010], specializing in incoporating side information with the factorization of the preference data. As mentioned in Section 2, GTagCDCF also builds itself on PMF, while specializing in incoporating side information from multiple domains with the factorization of the preference data.

In sum, GTagCDCF is designed to be effective for addressing challenges in CDCF to improve recommendation performance in different domains. In the remainder of the paper, we present our extensive experimental analysis that demonstrates the strength of GTagCDCF.

5. EXPERIMENTAL EVALUATION

In this section, we present the experiments we conducted to evaluate the proposed GTagCDCF. We first give a detailed description of the datasets that are used in our experiments, and the experimental setup. Then, we investigate the impact of the tradeoff parameters of GTagCDCF and the effectiveness of GTagCDCF. Finally, we demonstrate the performance of GTagCDCF by comparison with several CF and CDCF approaches. The research questions that need to be answered through the experiments are formulated as follows:

(1) Could GTagCDCF be effective for introducing mutual benefit between different recommender domains so as to improve recommendations in each domain?
(2) Could GTagCDCF outperform single-domain CF approaches?
(3) Could GTagCDCF outperform other state-of-the-art CDCF approaches by exploiting the common tags between domains?
(4) How would GTagCDCF perform for users with different rating profiles and different tagging behavior?
(5) How about the robustness of GTagCDCF in terms of the scale of datasets and the number of involved domains?

5.1. Datasets

Our experiments are organized as two parts: one for the two-domain case and the other for the three-domain case. For the two-domain case, the experiments are conducted on

two pairs of publicly available datasets. The first pair of datasets consists of one subset from MovieLens 10 million dataset[4] [Herlocker et al. 1999] and one subset from LibraryThing dataset[5] [Clements et al. 2010]. The original MovieLens dataset contains 10 million ratings from 71576 users and 10681 movies, and in the subset we select the first 5000 users and 5000 movies according to the identifiers in the original dataset. The original LibraryThing dataset contains ca. 750 thousand ratings from 7279 users and 37232 books, and in the subset we also select the first 5000 users and 5000 books. This subset selection was necessary, since some baselines are too computationally expensive to tackle larger datasets. Note that our choice of the subset selection procedure rather than random selection also ensures future experimental reproducibility. In the following, the two used subsets are denoted as ML1 (from MovieLens) and LT (from LibraryThing). Both ML1 and LT have 5-star rating scale with half star increments, representing a case that two different domains both have explicit ratings. In addition to ratings, the two domains also have 2277 common tags. Our target of experimenting on this pair of datasets is to investigate the effectiveness of our proposed GTagCDCF for benefitting different recommender domains that are both based on explicit ratings. For notation convenience, we refer to this pair of datasets as **P1** in the following. The second pair of datasets are from recent initiatives on information heterogeneity and fusion in recommender systems[6] [Cantador et al. 2011], where three datasets are provided with various types of user preferences and resources from three different domains, i.e., movie, webpage and music. Two datasets are used in our work. One dataset is an extended version of MovieLens 10 million dataset[7] with additional public information from IMDB[8] and Rotten Tomatoes[9] websites. In this following, we denote this dataset as ML2. ML2 contains ca. 850 thousand ratings (with the same scale of ML1 as mentioned before) from 2113 users and 10197 movies. The other dataset is collected from Last.fm[10] online music system, mainly consisting of user-artist listening information, i.e., the frequency that a user listened to songs from an music artist. In the following, we denote this dataset as LF. LF represents a dataset from implicit-feedback recommender systems, where no explicit user preferences are expressed. LF contains ca. 93 thousand user-artist listening relations, i.e., listening counts, from 1892 users and 17632 artists. In addition, there are in total 996 tags that are common to the two domains. Compared to **P1**, our target of experimenting on this pair of datasets is to investigate the effectiveness of our proposed GTagCDCF for benefitting different recommender domains that have different types of user preferences. We also refer to this pair of datasets as **P2** in the following of the paper. The statistics of all the datasets involved in the two-domain experiments are summarized in Table I.

For the three-domain case, the experiments are conducted by using the entire MovieLens 10 million dataset (denoted as ML-all, representing the movie domain,), the entire LibraryThing dataset (denoted as LT-all, representing the book domain) and the LF dataset (representing the music domain). Apart from the 2277 common tags between the ML-all dataset and the LT-all dataset as shown in Table I, there are 843 common tags between the ML-all dataset and the LF dataset, and 1152 common tags between the LT-all dataset and the LF dataset.

[4]http://www.grouplens.org/node/73
[5]http://dmirlab.tudelft.nl/users/maarten-clements
[6]http://ir.ii.uam.es/hetrec2011/datasets.html
[7]http://www.grouplens.org/node/462
[8]http://www.imdb.com
[9]http://www.rottentomatoes.com
[10]http://www.last.fm

Table I. Statistics of datasets in **P1** and **P2**.

	P1		**P2**	
	ML1	**LT**	**ML2**	**LF**
#users	5000	5000	2113	1892
#items	5000	5000	10109	17632
#preferences	584628	179419	855598	92834
sparseness	97.70%	99.30%	96%	99.70%
#common tags	2277		996	
#relations between users and common tags	559	157932	2551	9402
#relations between items and common tags	10778	97277	4016	28267

Note that in the following evaulation, our focus is on the two-domain case for the purpose of demonstrating the usefulness of the GTagCDCF model, while the experiments of the three-domain case only serve to validate the robustness of GTagCDCF to larger scale datasets and more than two domain use cases.

5.2. Experimental Setup

5.2.1. Experimental Protocol for Two-Domain Case. The experimental protocol we used in this work for the two-domain case is similar to some recent CDCF work [Li et al. 2009a; Li et al. 2009b]. In our experiments, we split each of all the datasets in both **P1** and **P2** into three sets containing different users, i.e., a training set, a validation set and a test set. For each dataset, the training set contains 60% randomly selected users and the users' preferences (ratings or listening counts) to items (movies, books or artists). The validation set contains 20% randomly selected users and their preferences to items. The test set contains the remaining 20% users and their preferences to items. The validation set is used to investigate the impact of different parameters in the proposed GTagCDCF algorithm and also tune parameters for all the baseline approaches. The test set is used to evaluate the performance of GTagCDCF and compare it with other baselines. For the purpose of investigating the performance of GTagCDCF for users with different rating profiles, we hold out preferences of users in the test set to be predicted, according to different conditions of user profile length (UPL). For example, under the condition of UPL=5, we use 5 randomly selected preferences for each user in the test set. The user profiles of length 5 are taken together with the training set as training data, and we use the remaining preferences of the users in the test set for evaluation. In our experiments, we evaluate three different conditions of UPL, i.e., UPL=5, 10, 15. For each condition of UPL, we generate 10 test data folds (each fold consisting of known preferences and holdout preferences) by random selection, and we report performance as an average across all of the 10 folds. In addition, in order to guarantee that each user in the test set has sufficient holdout preferences to be evaluated, we filter out users in the test set who have less than 20 preferences. Therefore, even for the condition of UPL=15, we still have at least 5 items for each user in the test set for evaluation.

5.2.2. Experimental Protocol for Three-Domain Case. As mentioned in Section 5.1, the purpose of the experiments for the three-domain case is only for validating the robustness of GTagCDCF. For this reason, we simply evaluate the performance of GTagCDCF by cross-validation without detailed examming different user profiles. Specifically, for each of the three datasets, we split it into 5 disjoint folds (i.e., each fold contains 20% preference data of the whole corresponding dataset.), from which we randomly select one for tuning the parameters in GTagCDCF and in the baseline approaches, and use the other four folds for cross-validation. The reported performance is averaged across all the four folds of each dataset.

5.2.3. Evaluation Metric. To be consistent with the majority of recent related work on CDCF [Li et al. 2009a; Li et al. 2009b; Pan et al. 2010; Pan et al. 2011], we adopt mean absolute error (MAE) as the evaluation metric for measuring recommendation performance on rating-based recommender domains, i.e., ML1 and LT in **P1** and ML2 in **P2**. Specifically, the definition of MAE is expressed as below:

$$MAE = \sum_{(i,j)\in T_E} |\hat{R}_{ij} - R_{ij}|/|T_E| \tag{19}$$

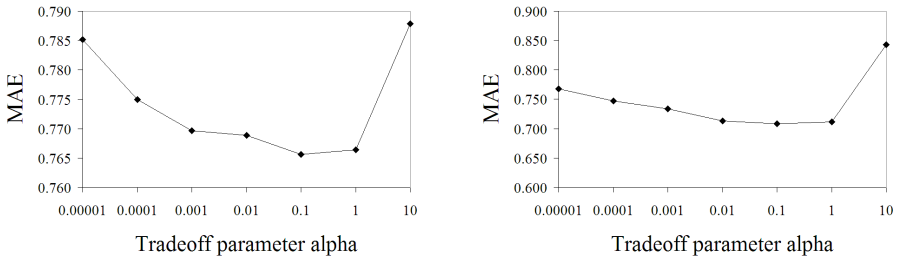
in which T_E denotes the set of user-item pairs whose ratings need to be predicted, and $|T_E|$ denotes the number of user-item pairs in the set. $\hat{R}_{ij}$ denotes the predicted rating for user i to item j, and R_{ij} the corresponding ground truth. A lower value MAE indicates a better recommendation performance.

For measuring recommendation performance on LF dataset, whose domain only has implicit user feedbacks, i.e., listening counts, we choose to use mean average precision (MAP) [Herlocker et al. 2004], a evaluation metric that is widely used to evaluate the quality of a ranking list. Specifically, MAP is defined as below:

$$MAP = \frac{1}{M_{te}} \sum_{i=1}^{M_{te}} \frac{\sum_{j=1}^{N_i}(rel_i(j) \times P_i@j)}{\sum_{j=1}^{N_i} rel_i(j)} \tag{20}$$

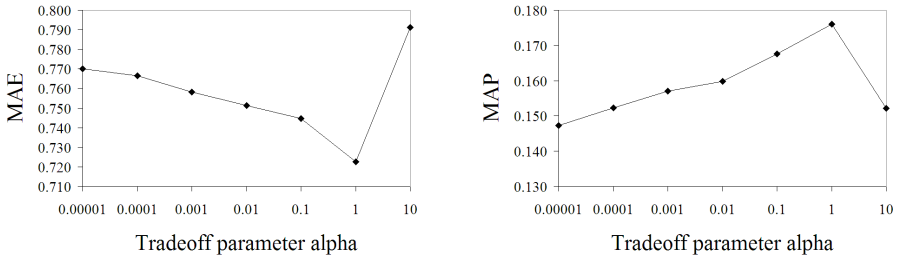
where M_{te} is the number of users for evaluation, and N_i denotes the number of recommended items (i.e., artists in LF dataset) for the user i. $rel_i(j)$ is a binary indicator, which is equal to 1 if the item of rank j is relevant to user i, and is equal to 0 otherwise. $P_i@j$ is the precision of the top j recommended items for the user i, i.e., the ratio of movies in the top j recommendation that are relevant to the user i. Higher values of MAP indicate a better recommendation performance. Note that in order to determine the relevance of each artist to a given user in the test set, we set a threshold of listening counts for each test user proportionate to the maximal listening count in her holdout set of feedbacks (listening counts to artists). The proportion coefficient for the threshold is set to 0.7 in our work. For example, if a holdout set of feedbacks for a test user has a maximal value as 1000, i.e., the maximal number of listening counts for that user to the artists that are to be ranked is 1000, then we regard the artists who are listened no less than 700 (i.e., $0.7\times$ 1000) times by the user as relevant ones. Although other strategies can be considered to determine the relevance of artists, they would not influence the result of our comparative study in the experiments as long as they lie in the assumption that the user would prefer one artist to the other if she listend to the one more than the other. In addition, note that our evaluation is only based on the artists who were listened by a given test user. The artists that were not listened by the user are not taken into account in the recommendation list, since there is no ground truth of user preference to those artists. Readers who have interest in this issue could further refer to the work of [Cremonesi et al. 2010; McLaughlin and Herlocker 2004]. For this reason, our evaluation on LF dataset is under a conservative manner, which could have underestimated the power of recommendation approaches, as shown later in section 5.5. However, this issue would not influence our comparative experimental evaluation in the following experiments, where we are interested in the relative strength of different approaches.

5.2.4. Parameter Setting. In the following experiments, we set the dimensionality of latent features in GTagCDCF to be 10. We notice that, just like in a common MF technique [Weimer et al. 2008], the performance of GTagCDCF did not substantially change when further increasing d, while more computational cost is required (cf. section 4.3). The regularization parameter λ of GTagCDCF is set to 0.01 for the two-domain case of **P1**, 0.001 for the two-domain case of **P2**, and 0.01 for the three-domain



(a) ML1 (b) LT

Fig. 4. The impact of tradeoff parameter α on the recommendation performance of GTagCDCF for the two domains in **P1**.



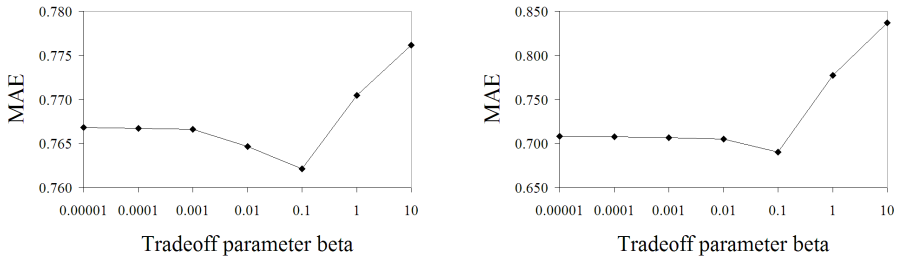
(a) ML2 (b) LF

Fig. 5. The impact of tradeoff parameter α on the recommendation performance of GTagCDCF for the two domains in **P2**.

case. The choice of λ for each case is based on our observation of the performance of GTagCDCF on the corresponding validation sets/folds. The impact of d and λ on recommendation performance has been widely investigated in related work particularly on MF [Koren et al. 2009; Salakhutdinov and Mnih 2008; Weimer et al. 2008]. In the following, we will focus on the impact of two tradeoff parameters α and β, which represent the key characteristics of GTagCDCF.

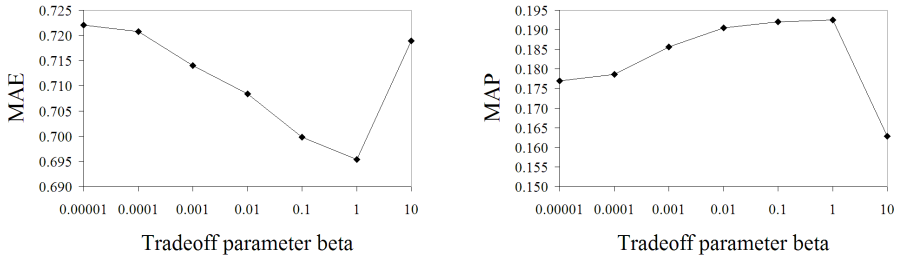
5.3. Impact of Tradeoff Parameters

The tradeoff parameters α and β in the proposed GTagCDCF algorithm influence the relative contributions from the cross-domain user-tag relations and the cross-domain item-tag relations. By using the validation set together with the training set in each of **P1** and **P2**, we investigate the impact of the tradeoff parameters on GTagCDCF by varying their values and measuring the recommendation performance in terms of MAE on ML1, LT and ML2 datasets and MAP on LF dataset. Note that in the experiments of this section, we set the condition of UPL=5 for the users in the validation set, while measuring recommendation performance according to the holdout set of ratings or implicit feedbacks as mentioned in the previous subsection. We first set $\beta = 0$ in Eq. (12) and investigate the impact of α, as shown in Fig. 4 and Fig. 5. It can be seen that for the cross-domain case of **P1**, the optimal value of α lies around 0.1 in terms of MAE for both ML1 domain and LT domain. In addition, it can also be seen that there is an optimal value of α, which lies around 1, in terms of MAE in ML2 domain and MAP in LF domain, when GTagCDCF is used for the case of **P2**. This observation indicates



(a) ML1 (b) LT

Fig. 6. The impact of tradeoff parameter β on the recommendation performance of GTagCDCF for the two domains in **P1**.



(a) ML2 (b) LF

Fig. 7. The impact of tradeoff parameter β on the recommendation performance of GTagCDCF for the two domains in **P2**.

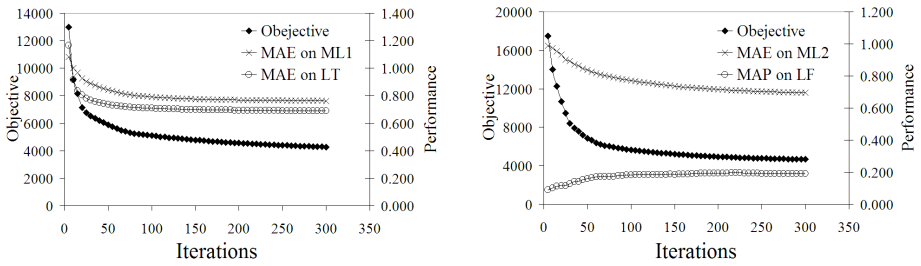
that by exploiting the common tags from different domains, GTagCDCF could benefit for improving latent user features in each domain, thus, resulting in an improved recommendation performance.

By adopting the optimal value of $\alpha = 0.1$ for the cross-domain case of **P1**, we further investigate the impact of tradeoff parameter β in Eq. (12) on the recommendation performance of GTagCDCF. We also further investigate the impact of tradeoff parameter β on GTagCDCF for the cross-domain case of **P2**, with α fixed to the optimal value 1. The results are shown in Fig. 6 and Fig. 7. As can be seen in Fig. 6, the optimal value of β lies around 0.1 in terms of MAE for both domains in **P1**. In addition, it can be also seen in Fig. 7 that the optimal value of β lies around 1 in terms of MAE in ML2 domain and MAP in LF domain. These results lead to an indication that by exploiting the common tags from different domains, GTagCDCF could also benefit for improving latent item features in each domain in addition to latent user features, thus, resulting in a further improved recommendation performance in each domain.

Based on our observations of the impact of tradeoff parameters as studied before, we could infer a positive answer to our first research question.

5.4. Effectiveness

In the second set of our experiments, we investigate the effectiveness of GTagCDCF, i.e., the effect of the minimization of the objective function in Eq. (12) on the improvement of the recommendation performance in different domains. Note that these experiments are also based on the valiation sets under the condition of UPL=5, and the



(a) P1 (b) P2

Fig. 8. The effectiveness of GTagCDCF in improving the recommendation performance (MAE or MAP) in different domains.

tradeoff parameters are adopted with the optimal values for **P1**and **P2**, respectively, as observed from the previous subsection. To this end, we demonstrate the variation of the output of the objective function and evaluation metrics, i.e., MAE for ML1, LT and ML2, and MAP for LF on the validation sets, simultaneously during the iterations of the optimization process. The results are shown in Fig. 8. As can be seen, when minimizing the objective function of GTagCDCF, the recommendation performance in each domain (for both **P1** and **P2**) improves along with algorithm iterations, nearly approaching to an convergence after 100 iterations. This observation allows us to confirm a positive answer to our first research question.

5.5. Performance Comparison in the Two-domain Case

In this subsection, we compare the performance of the proposed GTagCDCF with a set of alternative recommendation approaches listed below. As mentioned in section 5.2, the performance is reported based on the test set with 10 randomly separated folds. The tradeoff parameters are the optimal ones determined using the validation set as stated in section 5.3, i.e., $\alpha = 0.1$, $\beta = 0.1$ for ML1 and LT in **P1**, and $\alpha = 1$, $\beta = 1$ for ML2 and LF in **P2**.

— **UBCF**: User-based collaborative filtering [Herlocker et al. 1999] is used as a representative of memory-based CF approaches. The neighborhood size is tuned to 50 according to our observation of its performance on the validation sets.
— **PMF**: Probabilistic matrix factorization [Salakhutdinov and Mnih 2008] is a state-of-the-art model-based CF approach. The regularization parameter λ is tuned based on the validation set of each dataset, i.e., 0.01 for ML1 and LT, and 0.001 for ML2 and LF. Note that both UBCF and PMF are CF approaches for a single-domain use case.
— **CBT**: Codebook transfer [Li et al. 2009a] represents a state-of-the-art cross-domain CF approach. For the two domains in **P1**, we use one domain (e.g., ML1) as the auxiliary domain, which is used to construct a codebook, and the other domain (e.g., LT) as the target domain in which the recommendations are generated. Following the experimental protocol used in [Li et al. 2009a], we select 500 users and 500 items with most rating to construct the auxiliary domain, and set the number of clusters to 50 for both users and items. Note that CBT is not applicable for the case of **P2**, since it requires that two domains share a same rating scale.
— **RMGM**: Rating-matrix generative model [Li et al. 2009b] represents another state-of-the-art cross-domain CF approach. The implementation of RMGM is based on the

publicly available code package [11] supplied by the authors. As suggested in [Li et al. 2009b], we set the number of both the user and the item clusters to be 20. Similar to CBT, RMGM is not applicable for the case of **P2**, since it also requires that two domains share a same rating scale. Note that in both CBT and RMGM, the related parameters are also tuned based on the performance measured from the validation sets, as used in GTagCDCF.

— **TagCDCF**: Tag-induced cross-domain collaborative filtering is from our previous work [Shi et al. 2011], which represents a cross-domain CF approach with explicit knowledge between domains. The difference between TagCDCF and GTagCDCF has been discussed in section 4.5. Note that all the parameters involved in TagCDCF are tuned by the same means as in GTagCDCF according to observations from the validation sets.

The results of the comparative analysis are shown in Table II and Table III, from which we have several observations regarding the proposed algorithm GTagCDCF as summarized below.

First, as can be seen in both Table II and Table III, GTagCDCF significantly outperforms single-domain CF approaches, i.e., UBCF and PMF, to a large extent. Note that the significance of improvement is measured according to Wilcoxon signed rank significance test with $p<0.005$. For the two domains in **P1**, GTagCDCF improves over UBCF by 7∼8% on ML1 and 14∼19% on LT, and over PMF by up to 8.5% on ML1 and up to 10% on LT. Similar amount of improvement can also be observed in the cross-domain case of **P2**. In addition, we notice that other cross-domain CF approaches, i.e., CBT and RMGM, also consistently outperform single-domain approaches. In view of this result, we can conclude that GTagCDCF, as a CDCF approach, succeeds in benefiting recommendations in different domains by exploiting the common knowledge between domains. This observation allows us to confirm a positive answer to our second research question.

Second, we can observe from Table II that GTagCDCF also significantly outperforms CDCF approaches, i.e., by 3∼4% and ca. 3.5∼4.5% over CBT on ML1 and LT, and by i.e., by 1.5∼4.5% and ca. 3.5∼7.5% over RMGM on ML1 and LT. Note that TagCDCF also achieveed substantial improvements over CBT and RMGM. The results indicate that exploiting explicit common knowledge between domains could be more effective than relying on implicit common patterns between domains for the purpose of CDCF, verifying that GTagCDCF succeeds in introducing added value through common tags compared to existing CDCF approaches. In addition, GTagCDCF outperforms TagCDCF significatly as well. The relative improvement amounts to, for the case of **P1**, 1.4∼2.0% on ML1 and 0.5∼2.5% on LT in terms of MAE, and for the case of **P2**, 2.3∼5.6% on ML2 in terms of MAE and 6.8∼14.7% on LF in terms of MAP. This result indicates that GTagCDCF succeed in introducing the frequency of user-tag relations and the frequency of item-tag relations for maximally benefitting from common tags for improving recommendations in each domain. Those observations allow us to give a positive answer to our third research question.

Fourth, the improvement achieved by GTagCDCF is consistent for users with different profiles as shown in all the domains of both **P1** and **P2**. In addition, the results also demonstrate that the performance of GTagCDCF generally improves as the number of ratings in the user profile increases, indicating that GTagCDCF not only benefits from the common tags between domains, but also sufficiently exploiting rating data from both domains. Although most CF approaches have this characteristic, we emphasize that it is important to have recommender systems in a healthy cycle where users are

[11] http://sites.google.com/site/libin82cn/

Table II. Comparison of recommendation performance between GTagCDCF and the baseline approaches on datasets in **P1**.

	UPL=5		UPL=10		UPL=15	
	ML1 (MAE)	**LT** (MAE)	**ML1** (MAE)	**LT** (MAE)	**ML1** (MAE)	**LT** (MAE)
UBCF	0.833±0.009	0.857±0.009	0.785±0.004	0.795±0.005	0.766±0.003	0.771±0.004
PMF	0.831±0.010	0.771±0.009	0.771±0.010	0.761±0.006	0.753±0.010	0.769±0.010
CBT	0.792±0.009	0.729±0.010	0.752±0.003	0.694±0.008	0.737±0.003	0.682±0.006
RMGM	0.780±0.010	0.745±0.010	0.756±0.006	0.720±0.004	0.720±0.004	0.681±0.005
TagCDCF	0.777±0.008	0.709±0.004	0.735±0.004	0.674±0.003	0.719±0.003	0.663±0.002
GTagCDCF	**0.761***±0.008	**0.691***±0.004	**0.721***±0.004	**0.667***±0.002	**0.709***±0.003	**0.659***±0.003

Note: "*" denotes a significant improvement of GTagCDCF over all the other approaches, according to Wilcoxon signed rank significance test with $p<0.005$.

Table III. Comparison of recommendation performance between GTagCDCF and the baseline approaches on datasets in **P2**.

	UPL=5		UPL=10		UPL=15	
	ML2 (MAE)	**LF** (MAP)	**ML2** (MAE)	**LF** (MAP)	**ML2** (MAE)	**LF** (MAP)
UBCF	0.803±0.010	0.134±0.008	0.761±0.004	0.152±0.011	0.742±0.004	0.156±0.009
PMF	0.766±0.009	0.151±0.009	0.714±0.003	0.162±0.007	0.668±0.004	0.168±0.012
TagCDCF	0.748±0.009	0.162±0.010	0.686±0.003	0.175±0.011	0.653±0.003	0.190±0.013
GTagCDCF	**0.707***±0.007	**0.176***±0.008	**0.655***±0.002	**0.201***±0.008	**0.638***±0.003	**0.203***±0.013

Note: "*" denotes a significant improvement of GTagCDCF over all the other approaches, according to Wilcoxon signed rank significance test with $p<0.005$.

Table IV. Performance comparison for users with different number of used common tags in LT of **P1**.

# common tags (# users)	CBT	RMGM	TagCDCF	GTagCDCF
<5 (20)	0.736	0.716	0.728	0.712
5∼10 (62)	0.683	0.681	0.644	0.642
11∼20 (113)	0.671	0.678	0.659	0.652
20∼50 (193)	0.707	0.702	0.693	0.684
>50 (175)	0.665	0.664	0.641	0.639

encouraged to contribute more in order to attain better recommendations. The analysis positively answers our fourth question concerning the effect of user rating profiles on GTagCDCF.

Finally, we investigate the impact of the number of common tags, that serve as an explicit information source, on transferring knowledge between domains. This investigation is done by a closer look at the comparison between the CDCF approaches with tags, i.e., TagCDCF and GTagCDCF, and the CDCF approaches without tags, i.e., CBT and RMGM. We compare their performance on the users with different number of used common tags, as shown in Table IV. Note that we conduct this experiment only on the test set users of LT dataset under the condition of UPL=15, since the number of tags that the users assigned in ML1 varies in a much narrower range than LT so that it can hardly provide a convincing observation. As can be seen, for the users who used less than 5 common tags, the performance of TagCDCF and GTagCDCF is not obviously better (or even worse) than CBT or RMGM. However, users who used more than 5 common tags are able to benefit from TagCDCF and GTagCDCF, the performance of which becomes superior to CBT and RMGM. According to this observation, we can conclude that the common tags, serving as an explicit information source, could be more effective than implicit patterns mined from user preference for the purpose of transferring knowledge, and users who are active in tagging could particularly benefit from the proposed approach. We also notice that the number of users who used less than 5 common tags is not the majority, indicating the threshold of the used common tags for GTagCDCF to be beneficial is not difficult to reach in practice. This analysis provides our answer to the fourth research question concerning the user tagging behavior.

Table V. Comparison of recommendation performance between PMF, SVD++ and GTagCDCF on datasets in the three-domain case.

	ML-all (MAE)	**LT-all** (MAE)	**LF** (MAP)
PMF	0.610±0.007	0.656±0.002	0.416±0.005
SVD++	0.604±0.004	0.652±0.001	–
GTagCDCF	0.596^{*}±0.001	0.644^{*}±0.001	0.430^{*}±0.003

Note: "*" denotes a significant improvement of GTagCDCF over all the baselines, according to Wilcoxon signed rank significance test with $p<0.005$.

5.6. Performance Comparison in the Three-domain Case

As mentioned in Section 5.2, for validating the robustness of GTagCDCF to larger datasets and use cases with more than two domains, we evaluate GTagCDCF in the three-domain case. The tradeoff parameters are tuned to, $\alpha = 10$ and $\beta = 1$, based on the performance measured on the randomly selected fold, as described in Section 5.2. Then, the performance of GTagCDCF is measured via 4-fold cross-validation. Due to the relatively large size of the involved datasets, some baselines used in the two-domain case, as in the previous subsection, are too computational expensive to be deployed, such as UBCF and TagCDCF. In addition, the baselines CBT and RMGM in the previous subsection are not proposed for the use cases with more than two domains. For these reasons, only two single-domain baselines are employed for performance comparison. One is PMF, which is the same as described in the previous subsection. The other is SVD++ [Koren 2008], which is also a state-of-the-art CF approach renowned from the Netflix Prize Competition. The implementation of SVD++ in our experiments is based on the publicly available software MyMediaLite [Gantner et al. 2011]. The dimensionality of the latent factors in both of the baselines is set to 10, the same as for GTagCDCF. The related parameters in each of the baselines are also tuned based on the performance measured on the randomly selected fold, as described in Section 5.2.

The results are shown in Table V, from which we can observe that GTagCDCF still achieves significant improvement over PMF and SVD++ in each individual dataset, i.e., ca. 2% improvement on the ML-all and the LT-all datasets interms of MAE, and ca. 3% improvement on the LF dataset. Note that we did not find that SVD++ could generate reasonable result on the implicit feedback dataset, LF, thus, the corresponding entry was left empty. The observation from the experimental results indicates that GTagCDCF could be sufficiently robust for the use cases of large datasets and multiple domains, a positive answer to our last research question.

6. CONCLUSION AND FUTURE WORK

In this paper, we present a new approach, GTagCDCF, to address the cross-domain collaborative filtering problem. GTagCDCF exploits user-contributed tags as explicit links to connect users and items from different recommender domains, which allows for knowledge transfer between different domains in order to attain mutual benefit. GTagCDCF is designed as a generalized framework that incorporates all the relations among users, items and tags in different domains into a learning process, in which the latent features of users and items in each domain are enhanced by the knowledge transferred via the latent features of common tags from other domains. In addition, GTagCDCF is analyzed to have the advantage of being suitable for both homogeneous and heterogeneous cross-domain cases, and of being computationally scalable at the same time.

Based on two cross-domain cases, a homogeneous one and a heterogeneous one, our experimental results demonstrate that GTagCDCF could substantially outperform

several state-of-the-art single-domain and cross-domain CF approaches in both cases. This improvement is consistent for users with various rating profiles. We also show that GTagCDCF, as a generalized version, could substantially improve over TagCDCF, due to its ability to comprehensively utilize the information induced from common tags. In addition, we investigate the impact of user tagging behavior on the benefit that the users can attain from GTagCDCF and find that GTagCDCF could already become beneficial for those users who used just a few common tags, and it could be particularly beneficial for those who are active in tagging. Finally, we validate the robustness of GTagCDCF via the evaluation of its performance in the three-domain case.

For our future work, we would like to extend the GTagCDCF framework to exploit the common "context", rather than only tags, to further improve the mutual benefit between different recommender domains. Context has been recognized as a critical factor for recommender systems [Adomavicius et al. 2011]. Apart from exploiting contextual information for improving a particular recommender system, we would also use some types of contextual information that might be common to different recommender domains to establish cross-domain links. It would be interesting to investigate effective mechanisms to address the cross-domain collaborative filtering problem in the case of multiple types of cross-domain links.

REFERENCES

ADAMS, R. P., DAHL, G. E., AND MURRAY, I. 2010. Incorporating side information into probabilistic matrix factorization using Gaussian processes. In *Proceedings of the 26th Conference on Uncertainty in Artificial Intelligence*, P. Grünwald and P. Spirtes, Eds. 1–9.

ADOMAVICIUS, G., MOBASHER, B., RICCI, F., AND TUZHILIN, A. 2011. Context-aware recommender systems. *AI Magazine 32,* 3, 67–80.

ADOMAVICIUS, G. AND TUZHILIN, A. 2005. Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE Transactions on Kowledge and Data Engineering 17,* 6, 734–749.

BERKOVSKY, S., KUFLIK, T., AND RICCI, F. 2007. Cross-domain mediation in collaborative filtering. In *Proceedings of the 11th international conference on User Modeling*. UM '07. Springer-Verlag, Berlin, Heidelberg, 355–359.

BISHOP, C. M. 2006. *Pattern Recognition and Machine Learning*. Springer Berlin / Heidelberg.

BREESE, J. S., HECKERMAN, D., AND KADIE, C. M. 1998. Empirical analysis of predictive algorithms for collaborative filtering. In *UAI'98*. 43–52.

CACHEDA, F., CARNEIRO, V., FERNÁNDEZ, D., AND FORMOSO, V. 2011. Comparison of collaborative filtering algorithms: Limitations of current techniques and proposals for scalable, high-performance recommender systems. *ACM Trans. Web 5*, 2:1–2:33.

CANTADOR, I., BRUSILOVSKY, P., AND KUFLIK, T. 2011. 2nd workshop on information heterogeneity and fusion in recommender systems (hetrec 2011). In *Proceedings of the 5th ACM conference on Recommender systems*. RecSys 2011. ACM, New York, NY, USA.

CLEMENTS, M., DE VRIES, A. P., AND REINDERS, M. J. T. 2010. The influence of personalization on tag query length in social media search. *Inf. Process. Manage. 46*, 403–412.

CREMONESI, P., KOREN, Y., AND TURRIN, R. 2010. Performance of recommender algorithms on top-n recommendation tasks. In *Proceedings of the fourth ACM conference on Recommender systems*. RecSys '10. ACM, New York, NY, USA, 39–46.

DESHPANDE, M. AND KARYPIS, G. 2004. Item-based top-n recommendation algorithms. *ACM Trans. Inf. Syst. 22*, 143–177.

GANTNER, Z., RENDLE, S., FREUDENTHALER, C., AND SCHMIDT-THIEME, L. 2011. Mymedialite: a free recommender system library. In *Proceedings of the fifth ACM conference on Recommender systems*. RecSys '11. ACM, New York, NY, USA, 305–308.

HERLOCKER, J. L., KONSTAN, J. A., BORCHERS, A., AND RIEDL, J. 1999. An algorithmic framework for performing collaborative filtering. In *Proceedings of the 22nd annual international ACM SIGIR conference on Research and development in information retrieval*. SIGIR '99. ACM, New York, NY, USA, 230–237.

HERLOCKER, J. L., KONSTAN, J. A., TERVEEN, L. G., AND RIEDL, J. T. 2004. Evaluating collaborative filtering recommender systems. *ACM Trans. Inf. Syst. 22*, 5–53.

HOFMANN, T. 2004. Latent semantic models for collaborative filtering. *ACM Trans. Inf. Syst. 22*, 89–115.

KLEINBERG, J. AND SANDLER, M. 2004. Using mixture models for collaborative filtering. In *Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*. STOC '04. ACM, New York, NY, USA, 569–578.

KOLDA, T. G. AND SUN, J. 2008. Scalable tensor decompositions for multi-aspect data mining. In *Proceedings of the 2008 Eighth IEEE International Conference on Data Mining*. IEEE Computer Society, Washington, DC, USA, 363–372.

KOREN, Y. 2008. Factorization meets the neighborhood: a multifaceted collaborative filtering model. In *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*. KDD '08. ACM, New York, NY, USA, 426–434.

KOREN, Y., BELL, R., AND VOLINSKY, C. 2009. Matrix factorization techniques for recommender systems. *Computer 42*, 30–37.

LI, B., YANG, Q., AND XUE, X. 2009a. Can movies and books collaborate?: cross-domain collaborative filtering for sparsity reduction. In *Proceedings of the 21st international jont conference on Artifical intelligence*. IJCAI'09. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2052–2057.

LI, B., YANG, Q., AND XUE, X. 2009b. Transfer learning for collaborative filtering via a rating-matrix generative model. In *Proceedings of the 26th Annual International Conference on Machine Learning*. ICML '09. ACM, New York, NY, USA, 617–624.

LI, B., ZHU, X., LI, R., ZHANG, C., XUE, X., AND WU, X. 2011. Cross-domain collaborative filtering over time. In *Proceedings of Twenty-Second International Joint Conference on Artificial Intelligence*. AAAI, 2293–2298.

LIANG, H., XU, Y., LI, Y., NAYAK, R., AND TAO, X. 2010. Connecting users and items with weighted tags for personalized item recommendations. In *Proceedings of the 21st ACM conference on Hypertext and hypermedia*. HT '10. ACM, New York, NY, USA, 51–60.

LINDEN, G., SMITH, B., AND YORK, J. 2003. Amazon.com recommendations: Item-to-item collaborative filtering. *IEEE Internet Computing 7*, 76–80.

MA, H., KING, I., AND LYU, M. R. 2011. Learning to recommend with explicit and implicit social relations. *ACM Trans. Intell. Syst. Technol. 2*, 29:1–29:19.

MCLAUGHLIN, M. R. AND HERLOCKER, J. L. 2004. A collaborative filtering algorithm and evaluation metric that accurately model the user experience. In *Proceedings of the 27th annual international ACM SIGIR conference on Research and development in information retrieval*. SIGIR '04. ACM, New York, NY, USA, 329–336.

PAN, S. J. AND YANG, Q. 2010. A survey on transfer learning. *IEEE Trans. on Knowl. and Data Eng. 22*, 1345–1359.

PAN, W., LIU, N., XIANG, E., AND YANG, Q. 2011. Transfer learning to predict missing ratings via heterogeneous user feedbacks. In *Proceedings of the 22st international jont conference on Artifical intelligence*. IJCAI'11. Morgan Kaufmann Publishers Inc.

PAN, W., XIANG, E. W., LIU, N. N., AND YANG, Q. 2010. Transfer learning in collaborative filtering for sparsity reduction. In *AAAI'10*. 230–235.

PATEREK, A. 2007. Improving regularized singular value decomposition for collaborative filtering. In *Proc. KDD Cup Workshop at SIGKDD'07, 13th ACM Int. Conf. on Knowledge Discovery and Data Mining*. 39–42.

PORTEOUS, I., ASUNCION, A., AND WELLING, M. 2010. Bayesian matrix factorization with side information and dirichlet process mixtures. In *AAAI*. 563–568.

RENDLE, S. AND SCHMIDT-THIEME, L. 2010. Pairwise interaction tensor factorization for personalized tag recommendation. In *Proceedings of the third ACM international conference on Web search and data mining*. WSDM '10. ACM, New York, NY, USA, 81–90.

RESNICK, P., IACOVOU, N., SUCHAK, M., BERGSTROM, P., AND RIEDL, J. 1994. Grouplens: an open architecture for collaborative filtering of netnews. In *Proceedings of the 1994 ACM conference on Computer supported cooperative work*. CSCW '94. ACM, New York, NY, USA, 175–186.

ROBU, V., HALPIN, H., AND SHEPHERD, H. 2009. Emergence of consensus and shared vocabularies in collaborative tagging systems. *ACM Trans. Web 3*, 14:1–14:34.

SALAKHUTDINOV, R. AND MNIH, A. 2008. Probabilistic matrix factorization. In *Advances in Neural Information Processing Systems*. Vol. 20.

SARWAR, B., KARYPIS, G., KONSTAN, J., AND REIDL, J. 2001. Item-based collaborative filtering recommendation algorithms. In *Proceedings of the 10th international conference on World Wide Web*. WWW '01. ACM, New York, NY, USA, 285–295.

SEN, S., VIG, J., AND RIEDL, J. 2009. Tagommenders: connecting users to items through tags. In *Proceedings of the 18th international conference on World wide web*. WWW '09. ACM, New York, NY, USA, 671–680.

SHI, Y., LARSON, M., AND HANJALIC, A. 2011. Tags as bridges between domains: improving recommendation with tag-induced cross-domain collaborative filtering. In *Proceedings of the 19th international conference on User modeling, adaption, and personalization*. UMAP'11. Springer-Verlag, Berlin, Heidelberg, 305–316.

SHI, Y., LARSON, M., AND HANJALIC, A. 2013. Mining contextual movie similarity with matrix factorization for context-aware recommendation. *ACM Trans. Intell. Syst. Technol. 4,* 1.

SI, L. AND JIN, R. 2003. Flexible mixture model for collaborative filtering. In *ICML*. 704–711.

SINGH, A. P. AND GORDON, G. J. 2008. Relational learning via collective matrix factorization. In *Proceeding of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*. KDD '08. ACM, New York, NY, USA, 650–658.

SONG, Y., ZHANG, L., AND GILES, C. L. 2011. Automatic tag recommendation algorithms for social recommender systems. *ACM Trans. Web 5*, 4:1–4:31.

SYMEONIDIS, P., NANOPOULOS, A., AND MANOLOPOULOS, Y. 2010. A unified framework for providing recommendations in social tagging systems based on ternary semantic analysis. *IEEE Trans. on Knowl. and Data Eng. 22*, 179–192.

TANG, J., YAN, J., JI, L., ZHANG, M., GUO, S., LIU, N., WANG, X., AND CHEN, Z. 2011. Collaborative users' brand preference mining across multiple domains from implicit feedbacks. In *AAAI*. 477–482.

TIPPING, M. E. AND BISHOP, C. M. 1999. Probabilistic principal component analysis. *Journal of the Royal Statistical Society. Series B (Statistical Methodology) 61,* 3, 611–622.

TSO-SUTTER, K. H. L., MARINHO, L. B., AND SCHMIDT-THIE ME, L. 2008. Tag-aware recommender systems by fusion of collaborative filtering algorithms. In *Proceedings of the 2008 ACM symposium on Applied computing*. SAC '08. ACM, New York, NY, USA, 1995–1999.

UMYAROV, A. AND TUZHILIN, A. 2011. Using external aggregate ratings for improving individual recommendations. *ACM Trans. Web 5*, 3:1–3:40.

WANG, J., DE VRIES, A. P., AND REINDERS, M. J. T. 2006. Unifying user-based and item-based collaborative filtering approaches by similarity fusion. In *Proceedings of the 29th annual international ACM SIGIR conference on Research and development in information retrieval*. SIGIR '06. ACM, New York, NY, USA, 501–508.

WANG, X., SUN, J.-T., CHEN, Z., AND ZHAI, C. 2006. Latent semantic analysis for multiple-type interrelated data objects. In *Proceedings of the 29th annual international ACM SIGIR conference on Research and development in information retrieval*. SIGIR '06. ACM, New York, NY, USA, 236–243.

WEIMER, M., KARATZOGLOU, A., AND SMOLA, A. 2008. Improving maximum margin matrix factorization. *Mach. Learn. 72*, 263–276.

WETZKER, R., UMBRATH, W., AND SAID, A. 2009. A hybrid approach to item recommendation in folksonomies. In *Proceedings of the WSDM '09 Workshop on Exploiting Semantic Annotations in Information Retrieval*. ESAIR '09. ACM, New York, NY, USA, 25–29.

ZHANG, Y., CAO, B., AND YAN YEUNG, D. 2010. Multi-domain collaborative filtering. In *Proceedings of the 26th Conference on Uncertainty in Artificial Intelligence (UAI)*. UAI '10.

ZHEN, Y., LI, W.-J., AND YEUNG, D.-Y. 2009. Tagicofi: tag informed collaborative filtering. In *Proceedings of the third ACM conference on Recommender systems*. RecSys '09. ACM, New York, NY, USA, 69–76.

ZHU, S., YU, K., CHI, Y., AND GONG, Y. 2007. Combining content and link for classification using matrix factorization. In *Proceedings of the 30th annual international ACM SIGIR conference on Research and development in information retrieval*. SIGIR '07. ACM, New York, NY, USA, 487–494.